

# Convex Clustering: An Attractive Alternative to Hierarchical Clustering

Gary K. Chen<sup>1</sup>   Eric C. Chi<sup>2</sup>   John Michael O. Ranola<sup>3</sup>   Kenneth Lange<sup>4</sup>

<sup>1</sup>Department of Preventive Medicine (Biostatistics Division), University of Southern California, CA

<sup>2</sup>Department of Electrical and Computer Engineering, Rice University, TX

<sup>3</sup>Department of Statistics, University of Washington, WA

<sup>4</sup>Departments of Biomathematics, Human Genetics, and Statistics, University of Los Angeles, CA

## **Abstract**

The primary goal in cluster analysis is to discover natural groupings of objects. The field of cluster analysis is crowded with diverse methods that make special assumptions about data and address different scientific aims. Despite its shortcomings in accuracy, hierarchical clustering is the dominant clustering method in bioinformatics. Biologists find the trees constructed by hierarchical clustering visually appealing and in tune with their evolutionary perspective. Hierarchical clustering operates on multiple scales simultaneously. This is essential, for instance, in transcriptome data where one may be interested in making qualitative inferences about how lower-order relationships like gene modules lead to higher-order relationships like pathways or biological processes. The recently developed method of convex clustering preserves the visual appeal of hierarchical clustering while ameliorating its propensity to make false inferences in the presence of outliers and noise. The current paper exploits the proximal distance principle to construct a novel algorithm for solving the convex clustering problem. The solution paths generated by convex clustering reveal relationships between clusters that are hidden by static methods such as k-means clustering. Our convex clustering software separates parameters, accommodates missing data, and supports prior information on relationships. The software is implemented on ATI and nVidia graphics processing units (GPUs) for maximal speed. Several biological examples illustrate the strengths of convex clustering and the ability of the proximal distance algorithm to handle high-dimensional problems.

# 1 Author Summary

Pattern discovery is one of the most important goals of data-driven research. In the biological sciences hierarchical clustering has achieved a position of pre-eminence due to its ability to capture multiple levels of data granularity. Hierarchical clustering's visual displays of phylogenetic trees and gene-expression modules are indeed seductive. Despite its merits, hierarchical clustering is greedy by nature and often produces spurious clusters, particularly in the presence of substantial noise. This paper presents a relatively new alternative to hierarchical clustering known as convex clustering. Although convex clustering is more computationally demanding, it enjoys several advantages over hierarchical clustering and other traditional methods of clustering. Convex clustering delivers a uniquely defined clustering path that partially obviates the need for choosing an optimal number of clusters. Along the path small clusters gradually coalesce to form larger clusters. Clustering can be guided by external information through appropriately defined similarity weights. The current paper introduces a new algorithm for solving the convex clustering problem and applies it to a variety of biological datasets. Comparisons to hierarchical clustering demonstrate the superior robustness of convex clustering. Our genetics examples include inference of the demographic history of 52 populations across the world, a more detailed analysis of European demography, and a re-analysis of a well-known breast cancer expression dataset. Our new algorithm is a particular example of a class of MM algorithms known as proximal distance algorithms. The proximal distance convex clustering algorithm is inherently parallel and readily maps to modern many-core devices such as graphics processing units (GPUs). Our freely available software, `CONVEXCLUSTER`, exploits OpenCL routines that ensure compatibility across a variety of hardware environments.

# 2 Introduction

Pattern discovery is one of the primary goals of bioinformatics. Cluster analysis is a broad term for a variety of exploratory methods that reveal patterns based on similarities between data points. Well-known methods such as  $k$ -means invoke a fixed number of clusters. In complex biological data, the number of clusters is unknown in advance, and it is appealing to vary the number of clusters simultaneously with cluster assignment. Hierarchical clustering has been particularly helpful in understanding cluster granularity in gene-expression studies and other applications. In addition to producing easily visualized and interpretable results, hierarchical clustering is simple to implement and computationally quick. These are legitimate advantages, but they do not compensate for hierarchical clustering's instability to small data perturbations.

All principled methods of clustering attempt to decrease some criterion. Hierarchical clustering constructs a dendrogram by fusing or dividing observations (features). Fusion is referred to as agglomerative clustering and splitting as divisive clustering. Because of the greedy nature of the choices in hierarchical clustering, it returns clusters that are only locally optimal with respect to the underlying criterion. Solutions may vary depending on how the algorithm is initialized. To improve the chances of reaching a global minimum, multiple algorithm initializations must be tried. Even then there is no guarantee of optimality. A potentially greater handicap is that small perturbations in the data can lead to large changes in hierarchical clustering assignments. This propensity makes hierarchical clustering sensitive to outliers and promotes the formation of spurious clusters. In combination, the presence of local minima and the sensitivity to outliers lead to irreproducible results.

Although these objections are serious, a complete reformulation of hierarchical clustering is unnecessary. Recently Lindsten *et al.* [2011] and Hocking *et al.* [2011] introduced convex clustering based on minimizing a penalized sum of squares. The strict convexity and coercivity of their criterion guarantees a unique clustering path. The penalty term in convex clustering criterion accommodates prior information through nonuniform weights on data pairs. The solution paths of convex clustering retain the straightforward interpretability of hierarchical clustering while ameliorating its sensitivity to outliers and tendency to get trapped by local minima.

Despite the promise of convex clustering, there are two obstacles that stand in its way of becoming a practical tool in bioinformatics. The first is the challenge of large-scale problems. Current algorithms are computationally intensive and scale poorly on high-dimensional problems. A second obstacle is the minimal guidance currently available on how to choose penalty weights. Hocking *et al.* [2011] suggest some rules of thumb but offer little detailed advice. In our experience, the quality of the clustering path depends critically on well-designed weights. To address these issues, the current paper describes a fast new algorithm and a corresponding software implementation, CONVEXCLUSTER. Our advice on strategies for choosing penalty weights is grounded in a few practical biological examples. These examples support our conviction that convex clustering can be more nuanced than hierarchical clustering. Our examples include Fisher's Iris data from discriminant analysis, ethnicity clustering based on microsatellite genotypes from the Human Genome Diversity Project and SNP genotypes from the POPRES project, and breast cancer subtype classification via microarrays. In the POPRES data, we first reduce the genotypes to principal components and then use these to cluster. The paths computed under convex clustering expose features of the data hidden to less sophisticated clustering methods. The potential for understanding human evolution and history alone justify wider adoption of convex clustering.

### 3 Methods

Assume that there are  $n$  cases and  $p$  predictors. The more vivid language of graph theory speaks of nodes rather than cases and edges rather than pairs of cases. To implement convex clustering, Lindsten *et al.* [2011] suggest minimizing the penalized loss function

$$f_\mu(\mathbf{U}) = \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{u}_i - \mathbf{u}_j\| \quad (1)$$

relying on Euclidean norms. Here the column vector  $\mathbf{x}_i$  of the matrix  $\mathbf{X}$  records the predictors for case  $i$ , the column  $\mathbf{u}_i$  of the matrix  $\mathbf{U}$  denotes the cluster center assigned to case  $i$ ,  $\mu \geq 0$  tunes the strength of the penalty, and  $w_{ij} \geq 0$  weights the contribution of the case pair  $(i, j)$  to the penalty. The objective function  $f_\mu(\mathbf{U})$  treats the predictors symmetrically. If these range over widely varying scales, it is prudent to standardize each predictor to have mean 0 and variance 1.

Because the objective function  $f_\mu(\mathbf{U})$  is strictly convex and coercive, a unique minimum point exists for each value of  $\mu$ . When  $\mu = 0$ , the values  $\mathbf{u}_i = \mathbf{x}_i$  minimize  $f_\mu(\mathbf{U})$ , and there are as many clusters as cases. If the underlying graph is connected, then as  $\mu$  increases, cluster centers coalesce until all centers merge into a single cluster with all  $\mathbf{u}_i = \bar{\mathbf{x}}$ , the average of the data points  $\mathbf{x}_i$ . Although fission events as well as fusion events can in principle occur along the solution path, following the path as  $\mu$  increases typically reveals a hierarchical structure among the clusters. The weights encode prior information that guides clustering. Setting some of the weights equal to 0 reduces the computational load of minimizing  $f_\mu(\mathbf{U})$  in the proximal distance algorithm introduced next.

#### 3.1 The Proximal Distance Algorithm

The proximal distance principle is a new way of attacking constrained optimization problems [Lange and Keys, 2014]. The principle is capable of enforcing parsimony in parameter estimation while avoiding the shrinkage incurred by convex penalties such as the lasso. Shrinkage leads to imperfect model selection in addition to poor parameter estimates. The proximal distance principle seeks to minimize a function  $h(\mathbf{y})$ , possibly nonsmooth, subject to  $\mathbf{y} \in C$ , where  $C$  is a closed set, not necessarily convex. The set  $C$  encodes constraints such as sparsity. In the exact penalty method of Clarke [Borwein and Lewis, 2006; Clarke, 1990; Demyanov *et al.*, 2010], this constrained problem is replaced by the unconstrained problem of minimizing  $h(\mathbf{y}) + \rho \text{dist}(\mathbf{y}, C)$ , where  $\text{dist}(\mathbf{y}, C)$  denotes the Euclidean distance from  $\mathbf{y}$  to  $C$ . Note that  $\text{dist}(\mathbf{y}, C) = 0$  is a necessary and sufficient

condition for  $\mathbf{y} \in C$ . If  $\rho$  is chosen large enough, say bigger than a Lipschitz constant for  $h(\mathbf{y})$ , then the minima of the two problems coincide.

How does convex clustering fit in this abstract framework? Although the objective function  $f_\mu(\mathbf{U})$  is certainly nonsmooth, there are no constraints in sight. The strategy of parameter splitting introduces constraints to simplify the objective function. Since least squares problems are routine, the penalty terms constitute the intractable part of the objective function  $f_\mu(\mathbf{U})$ . One can simplify the term  $\|\mathbf{u}_i - \mathbf{u}_j\|$  by replacing the vector difference  $\mathbf{u}_i - \mathbf{u}_j$  by the single vector  $\mathbf{v}_{ij}$  and imposing the constraint  $\mathbf{v}_{ij} = \mathbf{u}_i - \mathbf{u}_j$ . Parameter splitting therefore leads to the revised objective function

$$g_\mu(\mathbf{U}, \mathbf{V}) = \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{v}_{ij}\| \quad (2)$$

with a simpler loss, an expanded set of parameters, and a linear constraint set  $C$  encapsulating the pairwise constraints  $\mathbf{v}_{ij} = \mathbf{u}_i - \mathbf{u}_j$ .

The proximal distance method undertakes minimization of  $h(\mathbf{y}) + \rho \text{dist}(\mathbf{y}, C)$  by a combination of approximation, the MM (majorization-minimization) principle [Borg and Groenen, 2005; Heiser, 1995; Hunter and Lange, 2004; Lange *et al.*, 2000; Wu and Lange, 2010], and an appeal to a combination of set projection [Deutsch, 2001] and proximal mapping [Parikh and Boyd, 2013]. The latter operations have been intensely studied for years and implemented in a host of special cases. Thus, the proximal distance principle encourages highly modular solutions to difficult optimization problems. Furthermore, most proximal distance algorithms benefit from parallelization.

Let us consider each of the ingredients of the proximal distance algorithm in turn, starting with approximation. The function  $\text{dist}(\mathbf{y}, C)$  is nonsmooth even when  $C$  is well behaved. For  $\epsilon > 0$  small, the revised distance  $\text{dist}_\epsilon(\mathbf{y}, C) = \sqrt{\text{dist}(\mathbf{y}, C)^2 + \epsilon}$  is differentiable and approximates  $\text{dist}(\mathbf{y}, C)$  well. The MM principle leads to algorithms that systematically decrease the objective function. In the case of minimizing  $f(\mathbf{y}) + \rho \text{dist}(\mathbf{y}, C)$  one can invoke the majorization  $\text{dist}(\mathbf{y}, C) \leq \|\mathbf{y} - P_C(\mathbf{y}_m)\|$ , where  $P_C(\mathbf{y}_m)$  is the projection of the current iterate  $\mathbf{y}_m$  onto the set  $C$ . By definition  $\text{dist}(\mathbf{y}_m, C) = \|\mathbf{y}_m - P_C(\mathbf{y}_m)\|$ , and  $P_C(\mathbf{y}_m)$  is a closest point in  $C$  to the point  $\mathbf{y}_m$ . For a closed nonconvex set, there may be multiple closest points; for a closed convex set there is exactly one.

According to the MM principle, minimizing the surrogate function

$$\frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{v}_{ij}\| + \rho \sqrt{\left\| \begin{pmatrix} \mathbf{U} \\ \mathbf{V} \end{pmatrix} - P_C \begin{pmatrix} \mathbf{U}_m \\ \mathbf{V}_m \end{pmatrix} \right\|^2} + \epsilon \quad (3)$$

drives the approximate objective function

$$\frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{v}_{ij}\| + \rho \sqrt{\text{dist} \left[ \begin{pmatrix} \mathbf{U} \\ \mathbf{V} \end{pmatrix}, C \right]^2} + \epsilon$$

downhill. The surrogate function (3) is still too complicated for our purposes. The remedy is another round of majorization. This time the majorization

$$\sqrt{t + \epsilon} \leq \sqrt{t_m + \epsilon} + \frac{1}{2\sqrt{t_m + \epsilon}}(t - t_m) \quad (4)$$

comes into play based on the concavity of the function  $\sqrt{t + \epsilon}$  for  $t \geq 0$ . As required by the MM principle, equality holds in the majorization (4) when  $t = t_m$ . Applying this majorization to the surrogate function (3) yields the new surrogate

$$\begin{aligned} h[(\mathbf{U}, \mathbf{V}) \mid (\mathbf{U}_m, \mathbf{V}_m)] &= \frac{1}{2} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{u}_i\|^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{v}_{ij}\| + \frac{\rho}{2d_m} \left\| \begin{pmatrix} \mathbf{U} \\ \mathbf{V} \end{pmatrix} - P_C \begin{pmatrix} \mathbf{U}_m \\ \mathbf{V}_m \end{pmatrix} \right\|^2 \\ d_m &= \sqrt{\left\| \begin{pmatrix} \mathbf{U}_m \\ \mathbf{V}_m \end{pmatrix} - P_C \begin{pmatrix} \mathbf{U}_m \\ \mathbf{V}_m \end{pmatrix} \right\|^2} + \epsilon \end{aligned} \quad (5)$$

up to an irrelevant constant. The surrogate function (5) resulting from these maneuvers separates all of the vectors  $\mathbf{u}_i$  and  $\mathbf{v}_{ij}$ . One can explicitly solve for the updates

$$\mathbf{u}_{n+1,i} = \frac{d_m}{d_m + \rho} \mathbf{x}_i + \frac{\rho}{d_m + \rho} \mathbf{a}_{ni},$$

where  $\mathbf{a}_{n,i}$  is the part of the projection pertaining to  $\mathbf{u}_i$ . The update of  $\mathbf{v}_{ij}$  involves shrinkage. Let  $\mathbf{b}_{n,ij}$  denote the part of the projection pertaining to  $\mathbf{v}_{ij}$ . Standard arguments from convex calculus [Lange, 2012] show that the minimum of  $\mu w_{ij} \|\mathbf{v}_{ij}\| + \frac{\rho}{2d_m} \|\mathbf{v}_{ij} - \mathbf{b}_{n,ij}\|^2$  is achieved by

$$\mathbf{v}_{n+1,ij} = \max \left\{ \left( 1 - \frac{\mu w_{ij} d_m}{\rho \|\mathbf{b}_{n,ij}\|} \right), 0 \right\} \mathbf{b}_{n,ij}.$$

In the exceptional case  $\mathbf{b}_{n,ij} = \mathbf{0}$ , the solution  $\mathbf{v}_{n+1,ij} = \mathbf{0}$  is clear from inspection of the  $\mathbf{v}_{ij}$  criterion. Both of these solution maps fall under the heading of proximal operators, hence, the name proximal distance algorithm.

If a weight  $w_{ij} = 0$ , then it is computationally foolish to introduce a difference vector  $\mathbf{v}_{ij}$ . In many applications, the weight matrix  $\mathbf{W} = (w_{ij})$  may be sparse. Given this lack of symmetry, one cannot expect to project analytically onto the constraint space. We now discuss a block descent algorithm for projection. Let  $E$  denote the set of edges  $\{i, j\}$  with positive weights  $w_{ij} = w_{ji}$ . Divide the neighborhood  $N_i$  of a node  $i$  into left and right neighborhoods  $L_i = \{j < i : w_{ji} > 0\}$  and  $R_i = \{j > i : w_{ij} > 0\}$ . Clearly  $N_i = L_i \cup R_i$ , and  $E = \cup_{i=1}^n N_i$ . Projection minimizes the criterion

$$\frac{1}{2} \sum_{i=1}^n \|\mathbf{u}_i - \tilde{\mathbf{u}}_i\|^2 + \frac{1}{2} \sum_{\{i,j\} \in E} \|\mathbf{u}_i - \mathbf{u}_j - \tilde{\mathbf{v}}_{ij}\|^2$$

for  $\tilde{\mathbf{U}}$  and  $\tilde{\mathbf{V}}$  given. It is unclear how to massage the stationarity equations

$$\mathbf{0} = \mathbf{u}_i - \tilde{\mathbf{u}}_i + \sum_{j \in R_i} (\mathbf{u}_i - \mathbf{u}_j - \tilde{\mathbf{v}}_{ij}) - \sum_{j \in L_i} (\mathbf{u}_j - \mathbf{u}_i - \tilde{\mathbf{v}}_{ji})$$

into a solvable form. However, the block updates

$$\mathbf{u}_i = \frac{1}{1 + |N_i|} \left( \tilde{\mathbf{u}}_i + \sum_{j \in R_i} \tilde{\mathbf{v}}_{ij} - \sum_{j \in L_i} \tilde{\mathbf{v}}_{ji} + \sum_{j \in N_i} \mathbf{u}_j \right)$$

are available. Here  $|N_i|$  denotes the cardinality of  $N_i$ . One cycle of the block descent algorithm updates  $\mathbf{u}_1$  through  $\mathbf{u}_n$  sequentially. This cycle is repeated until all of the vectors  $\mathbf{u}_i$  stabilize. Once convergence is achieved, one sets  $\mathbf{v}_{ij} = \mathbf{u}_i - \mathbf{u}_j$  for the relevant pairs.

### 3.2 Missing Data

The convex function (7) assumes no missing entries in the data matrix  $\mathbf{X}$ . It is straightforward to accommodate missing data by another round of majorization. Suppose  $\Gamma$  is the set of ordered index pairs  $(i, j)$  corresponding to the observed entries  $x_{ij}$  of  $\mathbf{X}$ . We now minimize the revised criterion

$$f_\mu(\mathbf{U}) = \frac{1}{2} \sum_{(i,j) \in \Gamma} (x_{ij} - u_{ij})^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{u}_i - \mathbf{u}_j\|, \quad (6)$$



which unfortunately lacks the symmetry of the original problem. To restore the lost symmetry, we invoke the majorization

$$\frac{1}{2} \sum_{(i,j) \in \Gamma} (x_{ij} - u_{ij})^2 \leq \frac{1}{2} \sum_{(i,j) \in \Gamma} (x_{ij} - u_{ij})^2 + \frac{1}{2} \sum_{(i,j) \notin \Gamma} (u_{mij} - u_{ij})^2,$$

where  $u_{mij}$  is a component of  $\mathbf{U}_m$ . In essence, the term  $(u_{mij} - u_{ij})^2$  majorizes 0. If the  $n \times p$  matrix  $\mathbf{Y} = (y_{ij})$  has entries  $y_{ij} = x_{ij}$  for  $(i, j) \in \Gamma$  and  $y_{ij} = u_{mij}$  for  $(i, j) \notin \Gamma$ , then in the minimization step of the proximal distance algorithm, we simply minimize the surrogate function

$$g_\mu(\mathbf{U}, \mathbf{V}) = \frac{1}{2} \sum_{i=1}^n \|\mathbf{y}_i - \mathbf{u}_i\|^2 + \mu \sum_{i < j} w_{ij} \|\mathbf{v}_{ij}\| \quad (7)$$

The rest of the proximal distance algorithm remains the same

### 3.3 Calibration of Weights

The pairwise weight  $w_{ij} = w_{ji}$  introduced in the penalty term of equation (1) determines the importance of similarity between nodes  $i$  and  $j$ . Two principles guide our choice of weights. First, the weight  $w_{ij}$  should be inversely proportional to the distance between the  $i$ th and  $j$ th points. This inverse relationship accords with intuition. As  $w_{ij}$  increases, the pressure for the  $i$ th and  $j$ th centroids to coalesce increases. If the weights  $w_{ij}$  are correlated with the similarity of the feature vectors  $\mathbf{x}_i$  and  $\mathbf{x}_j$ , then the pressure for their centroids to merge is especially great. Second, the weight matrix  $\mathbf{W}$  should be sparse. Despite the fact that small positive weights and zero weights lead to similar clustering paths, the computational advantages of zero weights cannot be ignored.

These observations prompt the following choice of weights. To maintain computational efficiency, it is helpful to focus on the  $k$  nearest neighbors of each node. We define the distance  $d_{ij}$  between two nodes  $i$  and  $j$  by the Euclidean norm  $\|\mathbf{x}_i - \mathbf{x}_j\|$  and write  $i \sim_k j$  if  $j$  occurs among the  $k$  nearest neighbors of  $i$  or vice versa. Based on these considerations the weights

$$w_{ij} = 1_{\{i \sim_k j\}} e^{-\phi d_{ij}^2} \quad (8)$$

are reasonable, where  $1_{\{i \sim_k j\}}$  is the indicator function of the event  $\{i \sim_k j\}$  and  $\phi \geq 0$  is a tuning constant. The case  $\phi = 0$  corresponds to uniform weights between nearest neighbors. When  $\phi$  is positive,  $w_{ij}$  strictly decreases as a function of  $d_{ij}$ . The relation  $i \sim_k j$  partitions the nodes into disjoint equivalence classes. Complete coalescence

of the nodes occurs as  $\mu$  increases if and only if there is a single equivalence class and the graph is connected. Using squared distances  $d_{ij}^2$  rather than distances  $d_{ij}$  induces more aggressive coalescence of nearby points and slower coalescence of distant points. In practice we normalize weights so that they sum to 1. This harmless tactic is equivalent to rescaling  $\mu$ .

This generic framework was proposed by Hocking *et al.* [2011]. We now discuss a strategy for leveraging additional information. When expert knowledge on the relationships among nodes is available and can be quantified, incorporating such knowledge may improve the clustering path. This must be done delicately so that prior information does not overwhelm observed data. In two of our examples, we integrate both genetic and geographic proximity measures in the weights  $w_{ij}$ . If  $\mathbf{x}_i$  and  $\mathbf{y}_i$  store the genotypes and GPS (global positioning system) coordinates on subject  $i$ , respectively, then the weighted average

$$d_{ij} = \alpha \|\mathbf{x}_i - \mathbf{x}_j\| + (1 - \alpha) \|\mathbf{y}_i - \mathbf{y}_j\|, \quad \alpha \in (0, 1), \quad (9)$$

serves as a composite distance helpful in clustering subjects. Observe that the components of the difference  $\mathbf{y}_i - \mathbf{y}_j$  must be computed in modulo arithmetic. Given a proper choice of the scaling constant  $\alpha$ , an even better alternative replaces  $\|\mathbf{y}_i - \mathbf{y}_j\|$  by the geodesic distance between  $i$  and  $j$ .

### 3.4 Evaluation of Clusters

Our program `CONVEXCLUSTER` minimizes the penalized loss (7) for a range of user specified  $\mu$  values. For each  $\mu$  the optimized matrix  $\mathbf{U}$  of cluster centers is stored in a temporary file for later construction of the cluster path. To facilitate visualization, `CONVEXCLUSTER` encourages users to project the cluster path onto any two principal components of the original data. The first example of Section 4 relies on the classical Iris data of discriminant analysis [Fisher, 1936]. This dataset contains 150 cases spread over three species. The Iris data can be downloaded from the UCI machine learning repository [Bache and Lichman, 2013]. For purposes of comparison, we also evaluated the clusters formed by agglomerative hierarchical clustering. Hierarchical clustering comes in several flavors; we chose UPGMA (Unweighted Pair Group Method with Arithmetic Mean) [Sokal and Michener, 1958] as implemented in the R function `hclust`. Although `hclust` offers six other options for merging clusters, UPGMA is probably the most reliable in reducing the detrimental effects of outliers since it averages information across all cluster members. UPGMA operates on a matrix of pairwise distances defined between nodes. In our genetics examples, we take these to be the distances defined by equation (9) and used in convex clustering.

## 4 Results

### 4.1 Impact of the Constants $k$ and $\phi$

To get a sense of the impact of the constants  $k$  and  $\phi$  on the Iris data, we generated cluster paths for various pairs  $(k, \phi)$ . As Figure 1 illustrates, the number of nearest neighbors  $k$  determines the connectivity of the underlying graph. Eventual coalescence only occurs for  $k = 50$ ; even then the apparent Iris-Versicolor outlier does not coalesce until very late. All values of  $k$  support a clear separation of Iris-Setosa from the other two species Iris-Versicolor and Iris-Virginica. Separation of Iris-Versicolor and Iris-Virginica into two different groups becomes discernible at  $k = 20$ . Subgroups within each species are evident for  $k = 5$  and  $k = 2$ . Improved resolution comes at a price; the two small two-member clusters seen in the top right corner of the main Iris-Versicolor cluster never fully coalesce with the main cluster when  $k = 2$ . The distance tuning constant  $\phi$  also exerts subtle influence along each row of Figure 1. This influence is more strongly felt for low values of  $k$ . Examination of the Iris data suggests exploring cluster granularity over a range of  $k$  values with  $\phi$  set to 0. One can find the minimum  $k$  ensuring full connectivity by combining bisection with either breadth-first search or depth-first search [Hopcroft and Tarjan, 1973]. Once the desired granularity is achieved,  $\phi$  can be increased to reveal more subtle details. Note that increasing  $\phi$  sends most weights between  $k$  nearest neighbors to 0. As previously noted, the proximal distance algorithm takes substantially more iterations to converge for large values of  $\phi$ .

### 4.2 Cluster Accuracy in the Presence of Noise

Although agglomerative hierarchical clustering is computationally efficient, it tends to create spurious clusters due to its greedy nature. In particular, it can falter in the face of noisy data. To test this hypothesis, we simulated new data from the Iris data. In creating a dataset, we perturbed each row of the data matrix  $X$  by adding normal deviates with mean 0 and standard deviation equal to the sample standard deviation  $s^2$  of the corresponding predictor multiplied by a constant  $c$ . We then clustered the data points into three clusters and counted the number of inconsistencies between cluster labels and species labels. For convex clustering, visual inspection of the converged clustering paths reveals roughly three major clusters for values of  $k$  between 5 and 15. With hierarchical clustering, three clusters were constructed by choosing a cut point on the full tree intersecting three branches. Both methods performed equally well on the original dataset, misclassifying the same 14 Iris-Virginica specimens as Iris-Versicolor for an overall error rate of  $14/150 = 0.093$ . One of these errors is possibly a mis-attribution of species; the remaining may represent hybrid plants. Table 1 summarizes error rates averaged over 100 replicates under the two methods. Examination of the table suggests that convex clustering is indeed

more accurate in the face of noise over a wide range of  $k$  values.

### 4.3 Cluster Accuracy with Missing Values

We carried out a second simulation study on the Iris data to assess the impact of varying levels of missingness on cluster inference. Because the Iris data includes only four features, simply selecting entries of the data matrix at random can lead to cases retaining no data. To avoid these degeneracies, we randomly selected cases and then a random feature from each case for deletion. Given cases rates of 25%, 50%, 75%, and 100%, the proportion of missing observations consequently ranged from 5% to 25%. Hierarchical clustering with missing data requires that either cases with missing entries be omitted or that missing entries be imputed. We employed the second strategy, filling in missing entries by multiple imputation as implemented in the R package `MI` [Su *et al.*, 2011]. Hierarchical clustering was then applied to the completed data. For convex clustering, we also applied multiple imputation, but for the sole purpose of computing the convex clustering weights. We then applied convex clustering to the original incomplete data under the objective function (6). Accuracy for each method was estimated in the same manner as the previous simulations. The error rates in Table 2 suggest that convex clustering does indeed outperform hierarchical clustering in the presence of missing data.

### 4.4 Inference of Ethnicity

As genotyping costs have dropped in recent years, it has become straightforward to relate ethnicity to subtle genetic variations. Several software tools are now available for this purpose. For example, the programs `STRUCTURE` [Pritchard *et al.*, 2000] and `ADMIXTURE` [Alexander *et al.*, 2009] estimate a subject's admixture proportions across a set of predefined or inferred ancestral populations. `EIGENSTRAT` [Price *et al.*, 2006] employs a handful of principal components to explain ethnic variation. Principal component analysis (PCA) is attractive due to its speed and ease of visualization. Clustering can also separate people by ethnicity if individuals of mixed ethnicity are omitted. The advantage of convex clustering is that one can follow the dynamic behavior of the relationship clusters along the regularization path.

#### 4.4.1 World-wide Genetic Diversity

For a practical demonstration of convex clustering, we now turn to the Human Genome Diversity Project (HGDP). This collaboration makes several datasets publicly available that vary in marker type (SNPs versus microsatellites) and sample size. The HGDP 2002 dataset considered here includes 1,056 individuals from 52

populations genotyped at 377 autosomal microsatellites [Rosenberg *et al.*, 2002]. Care must be taken in analyzing microsatellites since, in contrast to SNPs, they display more alleles and greater levels of polymorphism. Recall that an allele at a microsatellite approximates the number of short tandem repeats of some simple motif. Because treating microsatellite genotypes as continuous variables is problematic, we encode each microsatellite genotype as a sequence of allele counts. Each count ranges from 0 to 2, and there are as many count variables as alleles. This encoding yields a revised 2002 dataset with the 377 microsatellite genotypes expanded to 4,682 different attributes.

As expected, these data exhibit clines in allele frequencies [Kittles and Weiss, 2003]. To take advantage of the correlation between geographic separation and ethnic similarity, we defined penalty weights  $w_{ij}$  according to the composite distance in equation (9) with constant  $\alpha = 0.5$ . Figure 2 plots cluster paths for these data given the settings  $\phi = 1$  and  $k = 4$ . With  $k = 4$  nearest neighbors, we observe broad-scale clustering events that link up the major continental groups. In the north, Europeans fall into a single cluster, later joined by populations from the Middle East. In the east the Chinese merge into a cluster that subsequently merges with two Oceania populations from New Guinea. This mega cluster then merges with various Central Asian populations of predominantly Pakistani origin. In the west five Central/South American populations cluster, and in the south six African populations cluster. Considering the continental clusters, the first two to coalesce are the American cluster and the Central/East Asian cluster. This accords with known links between East Asians and American Indians, who crossed the Bering strait, possibly multiple times, during the Ice Age. Figure 3 depicts finer grained events exposed by setting  $k = 1$ . Along the western axis, taking  $k = 1$  is uninformative, but among the African populations along the southern axis, we observe three major clusters: a two-member cluster representing the two Pygmy sub-groups; a three-member cluster comprising Bantu-speaking peoples from Kenya, Yorubans from Nigeria, and Mandenkas from Senegal; and finally a singleton cluster for the San from Namibia. These results are consistent with a recent phylogenetic study [Li *et al.*, 2008] that found the San to be the most isolated of the African populations, followed by the two Pygmy populations, and finally the three Bantu-language populations. Along the eastern axis, the two Papua New Guinea populations cluster together and do not join the remaining Asian populations.

Figures 4 and 5 focus on related populations along the eastern and northern axes of East Asia, respectively. Most of the Chinese populations along the eastern axis appear to coalesce simultaneously. Some of the other populations along the northern border of China coalesce early. The Hezhen and Oroqen peoples reside predominantly in the Heilongjiang province of northeast China [census bureau, 1990, 2000]. These two populations cluster early with the inner Mongolians and the Xibo population, who occupy northeast China and the northwest region of

Xinjiang province. Three distinct clusters of Middle Easterners, Central Asians, and Europeans occur along the northern axis. All European populations except for the Russian populations are grouped into a single cluster. The two Russian populations instead merge with a second cluster that includes three populations from Israel. The Mozabites, who coalesce late with this cluster, exhibit high frequencies of North African haplotypes as previously noted in the literature [Rosenberg *et al.*, 2006; Coudray *et al.*, 2009]. A third cluster within Central Asia unite Pakistani populations with Uygurs from China. Within this cluster, the Brahui, Balochi, and Makran populations of the Baluchistan province of northwestern Pakistan coalesce early with the Sindhi people of the Sindh province on the eastern border of Baluchistan. Later coalescing populations include the Hazara, Uygurs, and Kalash. The Hazaras of Pakistan and the Uygurs of China share common Mongolian and Turkic ancestry and some physical attributes [Qamar *et al.*, 2002; Ablimit *et al.*, 2013]. A previous admixture analysis using high-density SNP data supports our observation that the Kalash people constitute a single distinct cluster, one of seven clusters separating all of the populations covered in the HGDP data [Rosenberg *et al.*, 2006].

The dendrogram in Figure 6 presents the output of hierarchical clustering. Results are largely consistent across the two methods. For example, in East Asia both methods cluster the Chinese groups in the north of China with the Yakut of Siberia and the Japanese. The two populations from Papua New Guinea form their own cluster in both methods. However, some key differences occur. Convex clustering infers a closer genetic similarity between the Uygurs and the Hazara than suggested by hierarchical clustering. Based on evidence from admixture analyses [Rosenberg *et al.*, 2006] and convex clustering, one expects the Kalash to coalesce after the Uygur and other Central South Asian populations have coalesced. Although the non-Russian populations from Europe clustered in one group under both methods, hierarchical clustering included the Mozabites within this group as well, suggesting that they are more genetically related to Italians from Bergamo than Italians from Tuscany. This observation is inconsistent with our findings and previous work [Coudray *et al.*, 2009; Rosenberg *et al.*, 2006]. Perhaps, the most puzzling discrepancy is the case where hierarchical clustering coalesces the African San population very late, only after all populations outside the Americas have coalesced.

#### 4.4.2 Population Structure of Europe

We next investigate whether convex clustering can glean further insights into the population structure of Europe. The POPRES resource archives high-density genotypes generated on the Illumina 550k microarray platform [Nelson *et al.*, 2008]. Version 2 of POPRES contains genotype and phenotype data on 4,077 subjects genotyped across 457,297 SNPs. For this analysis, we include only non-admixed Europeans who report all four grandparents of the same ethnicity. This leaves 1,896 subjects. SNP data presents advantages and disadvantages

compared to microsatellite data. Dense marker panels may be more sensitive to subtle differences driven by population events such as migration, expansion, and bottlenecks. Challenges include the lower information content of biallelic markers and the correlations between markers caused by linkage disequilibrium (LD). After considerable experimentation, we found that the leading principal components offered more insight into population structure than the raw genotypes themselves. We employed EIGENSTRAT to extract the ten leading principal components from the genotype matrix. EIGENSTRAT prunes SNPs in LD with  $r^2$  exceeding a user-specified threshold [Price *et al.*, 2006]. In our case the threshold 0.8 discards all but 276,823 nearly independent SNPs. Our choice of the composite distance defined in equation (9) places equal weight ( $\alpha = 0.5$ ) on genetic distances and GPS distances between the capital cities of participants. To ease visualization, our figures display a maximum of 20 subjects from each ethnicity, for a total of 370 subjects. The computed convex clustering path is projected onto the first two principal components of the POPRES data; these components capture geographic east-west and north-south axes, respectively.

In the Iris and the HGDP datasets, the number of nearest neighbors  $k$  was more critical in resolving cluster evolution than the tuning constant  $\phi$ . In the European POPRES data, where inter-class differences are more subtle, increasing  $\phi$  can be critical in resolving details for  $k$  large. Figure 7 depicts a clustering path with  $k = 40$  neighbors and  $\phi = 0$ . Increasing  $\phi$  to 10 gives a similar clustering pattern, except that each of the major trunks coalesce before converging to the origin. Thus, Figure 8 shows several major clusters connected by five major trunks. Spain and Portugal constitute a major cluster in the southwest trunk. The southeast trunk includes Italy and southeast Europe; these populations eventually merge into a single cluster. The northeast trunk defines a cluster that includes Poland, Russia, Ukraine, the Czech Republic, Hungary, and Slovenia. Norway, Sweden, and Germany cluster along the northern trunk, and the British Isles merge with Belgium and the Netherlands to form the northwest trunk. A large cluster comprising France and the Swiss linguistic groups (French, German, and Italian) constitute the western trunk. Replotting the clustering path with  $\phi = 1$  and  $k = 3$  shows Norway and Sweden breaking away from Germany and forming their own disjoint cluster. France breaks away from the Swiss groups to form its own disjoint cluster. Along the south trunk, Italy now separates from southeast Europe and eventually clusters with the Swiss-Italians.

Figure 9 depicts the clustering path of southeast Europe, where West Slavic languages predominate. Here Greece first coalesces with Macedonia, a Slavic population bordering Greece on the north. A cluster comprising Bosnia-Herzegovina and Serbia merges with Romania, before merging into the primary trunk of southeast Europe. Finally at the northern end of the trunk, a cluster formed by Croatia and Slovenia form its own cluster. The groups in the Bosnia-Herzegovina cluster and the Macedonian cluster are consistent with local history. Poland

and Russia cluster in the northern most branch of the northeast trunk (Figure 10). The Czech-Republic, Austria, and Hungary define a distinct cluster along the southern branch. Given that Austria conquered Hungary in 1699 and established rule over Bohemia (the predecessor to modern Czechs) as early as 1526, these results are not surprising.

In the POPRES data, convex clustering and hierarchical clustering occasionally disagree. For example, hierarchical clustering merges the Netherlands and Belgium with Britain before it merges Britain with Ireland and Scotland (Figure 11). In light of the geography and history of Britain, it is reasonable to expect Britain to first merge with Scotland and Ireland. Convex clustering produces this intuitive cluster, which eventually merges with the neighboring cluster comprising Belgium and the Netherlands. Owing to a few outliers, the greedy nature of hierarchical clustering appears to force a spurious coalescence, which cannot be repaired until later. Another discrepancy occurs in clustering the Swiss linguistic groups. Convex clustering first groups the Swiss-German, Swiss-French, and Swiss-Italian into a single Swiss cluster. Hierarchical clustering groups France with this cluster. At the next higher level, hierarchical clustering fails to cluster Italy with the Swiss and instead merges it with Greece and populations from the former Yugoslavia. Convex clustering, on the other hand, produces a more intuitive path, merging Italy with the Swiss before joining both to the southeast European trunk.

## 4.5 Inferring Cancer Subtypes

It is well accepted that cancers of a given tissue often fall into different subtypes. In breast cancer for instance, patients with tumors that are estrogen receptor (ER) negative are less responsive to hormone based treatment than those possessing active estrogen receptors (ER) [Roche *et al.*, 2003]. High-throughput platforms such as gene-expression microarrays and RNA-Seq have enabled researchers to classify cancer patients based on their molecular phenotypes. Hierarchical clustering by Perou *et al.* [2000] established five gene-expression profiles across 9216 genes in 84 breast-cancer patients. Among the 84 patients, only 16 also had a clinical assessment of ER status. It is natural to ask whether convex clustering can discriminate between the two ER tumor types. Under the tuning constants  $\phi = .5$  and  $k = 1$ , convex clustering recovers two distinct clusters. Figure 12 projects the cluster centers along the cluster path on the first and third principal components of the original data. The left and right clusters correspond roughly to ER positive and ER negative tumors, respectively. Two ER negative tumors appear as outliers in the ER positive cluster. Within the ER positive cluster, two ER positive samples with ER-B2 positive status cluster before merging with the remaining samples. Figure 13 depicts an unrooted tree derived from hierarchical clustering. In comparison to convex clustering, the cluster definitions are vague, with most clusters containing two or three members. This could be an artifact of the hard binary choices imposed by



hierarchical clustering. The two ER-B2 positive samples that clustered together in convex clustering appear in distant clusters under hierarchical clustering.

## 4.6 Run-time benchmarks

For a dataset with a large number of attributes, parallelization can substantially reduce run times. CONVEXCLUSTER includes code written in OpenCL, a language designed to run on many-core devices such as GPUs. For each of the three genetic analyses presented above, we recorded the total run-time along the entire regularization path using standard C++ code for the CPU and OpenCL code for the GPU. For the sake of comparison, we also recorded run-times for CLUSTERPATH on the same datasets and weighting schemes. Table 3 records the average run time to minimize the objective function averaged over all values of the regularization parameter. We chose this strategy because CLUSTERPATH does not allow users to pre-specify a grid of regularization values. The bottom line is that CONVEXCLUSTER required only 16%, 47%, and 75% of the time required by CLUSTERPATH to fit the HGDP, POPRES, and breast cancer datasets respectively. When a GPU is available, further improvements can potentially be realized. On an nVidia C2050 GPU, CONVEXCLUSTER enjoys speed improvements of 4.6 and 5.5 fold over the CPU version for the HGDP and breast cancer examples. In contrast, on the POPRES example, the GPU version is actually 3.5 fold slower than the CPU version. GPU programs suffer when a significant amount of time is spent transferring data to and from the GPUs. In its current form, CONVEXCLUSTER reads the updated matrix  $U$  from the GPUs at each point on the  $\mu$ -regularization path before saving the data to disk. This large I/O overhead can overwhelm gains from parallelization for low-dimensional datasets such as the POPRES data. In general, GPU implementations of standard algorithms require a high degree of parallelization, limited data transfers between the master CPU and the slave GPUs, and maximal synchrony of the GPUs. Depending on the nature of the clustering data, CONVEXCLUSTER satisfies these requirements.

## 5 Discussion

The literature on cluster analysis is enormous. Each clustering method has advantages in either simplicity, speed, reliability, interpretability, or scalability. If the number of clusters is known in advance, then  $k$ -means clustering is usually preferred. In convex clustering one can often achieve a predetermined number of clusters by varying the number of nearest neighbors and following the solution path to its final destination. Alternatively, if the underlying graph is fully connected, then one can follow the solution path until  $k$  clusters appear. The downside of  $k$ -means clustering is that it offers no insight into cluster similarity. If the goal in clustering is to

obtain a snapshot of the relationships among observed data points at different levels of granularity, the choices are limited, and most biologists opt for hierarchical clustering. Hierarchical clustering is notable for its speed and visual appeal. Balanced against these assets is its sensitivity to poor starting values and outliers. Our perturbations of the Iris data demonstrate the latter weakness. Convex clustering occupies an enviable middle ground between  $k$ -means clustering and hierarchical clustering. Our extensive exploration of the HGDP and POPRES datasets showcase the subtle solutions paths of convex clustering. These paths offer considerable insights into population history and correct some of the greedy mistakes of hierarchical clustering.

Given the novelty of convex clustering [Lindsten *et al.*, 2011], it is hardly surprising that only a single previous program, CLUSTERPATH, implements it [Hocking *et al.*, 2011]. CONVEXCLUSTER and CLUSTERPATH perform similarly on modest problems such as the Iris data. Unfortunately, on large datasets such as the HGDP data, CLUSTERPATH depletes all available memory and fails. Furthermore, CLUSTERPATH lacks two features that work to the advantage of convex clustering. First, it does not support disconnected graphs defined by sparse weights. In our breast cancer example, clustering with disconnected graphs reveals fine-grained details. Second, CLUSTERPATH does not allow for missing entries in the data matrix. The current paper documents CONVEXCLUSTER's ability to scale realistically to dimensions typical of modern genomic data. A combination of careful algorithmic development and exploitation of modern many-core chipsets lies behind CONVEXCLUSTER. The proximal distance algorithm propelling CONVEXCLUSTER separates parameters and enables massive parallelization. OpenCL made it relatively easy to implement parallel versions of our original serial code. Further speedups are possible. For instance, CONVEXCLUSTER spends an inordinate amount of execution time moving matrices over relatively slow I/O channels in preparation for plotting. One could easily project the data to principal components on each GPU itself prior to data transfer. More recent ATI or nVidia GPUs should improve the speedups on high-dimensional data mentioned here.

Convex clustering also shows promise as a building block for more sophisticated exploratory tools in computational biology. In a companion paper Chi *et al.* [2014] introduce a convex formulation of the biclustering problem. In biclustering one seeks to cluster both observations and features simultaneously in a data matrix. Cancer subtype discovery can be formulated as a biclustering problem in which gene expression data is partitioned into a checkerboard-like pattern highlighting the associations between groups of patients and the groups of genes that distinguish them. To bicluster a data matrix, hierarchical clustering can be applied independently to the rows and columns of the matrix. Convex biclustering produces more stable biclusterings while retaining the interpretability of hierarchical biclustering. Convex biclustering requires repeatedly solving convex clustering subproblems.

The field of cluster analysis is crowded with so many competing methods that it would foolish to conclude

that convex clustering is uniformly superior. Our goal of illustrating the versatility of convex clustering is more modest. The reflex reaction of most biologists is to employ hierarchical or  $k$ -means clustering. We suggest that biologists take a second look. Convex clustering's ability to reliably deliver an entire solution path is compelling. The insights discussed here will enhance the careful exploration of many big datasets. The present algorithm, and indeed the present formulation of convex clustering, are unlikely to be the last words on the subject. We encourage other computational biologists and statisticians to refine these promising tools. CONVEXCLUSTER can be freely downloaded from the UCLA Human Genetics web site at <http://www.genetics.ucla.edu/software/> for analysis and comparison purposes.

## References

- Ablimit, A., Qin, W., Shan, W., Wu, W., Ling, F., Ling, K. H., Zhao, C., Zhang, F., Ma, Z., and Zheng, X. (2013). Genetic diversities of cytochrome B in Xinjiang Uyghur unveiled its origin and migration history. *BMC Genet.*, **14**, 100.
- Alexander, D. H., Novembre, J., and Lange, K. (2009). Fast model-based estimation of ancestry in unrelated individuals. *Genome research*, **19**(9), 1655–1664.
- Bache, K. and Lichman, M. (2013). UCI machine learning repository.
- Borg, I. and Groenen, P. J. (2005). *Modern Multidimensional Scaling: Theory and Applications*. Springer.
- Borwein, J. M. and Lewis, A. S. (2006). *Convex Analysis and Nonlinear Optimization*, volume 3 of *CMS Books in Mathematics*. Springer, second edition.
- census bureau, C. (1990). The fourth population census of china in 1990.
- census bureau, C. (2000). Population census of china in 2000.
- Chi, E. C., Allen, G. I., and Baraniuk, R. G. (2014). Convex biclustering. arXiv:1408.0856 [stat.ME].
- Clarke, F. H. (1990). *Optimization and Nonsmooth Analysis*, volume 5 of *Classics in Applied Mathematics*. SIAM.
- Coudray, C., Olivieri, A., Achilli, A., Pala, M., Melhaoui, M., Cherkaoui, M., El-Chennawi, F., Kossmann, M., Torroni, A., and Dugoujon, J. M. (2009). The complex and diversified mitochondrial gene pool of Berber populations. *Ann. Hum. Genet.*, **73**(2), 196–214.
- Demyanov, V. F., Fletcher, R., Terlaky, T., Di Pillo, G., and Schoen, F. (2010). *Nonlinear Optimization*. Springer.
- Deutsch, F. (2001). *Best Approximation in Inner Product Spaces*, volume 7 of *CMS Books in Mathematics*. Springer-Verlag.
- Fisher, R. A. (1936). The use of multiple measurements in taxonomic problems. *Annals of eugenics*, **7**(2), 179–188.
- Heiser, W. J. (1995). Convergent computation by iterative majorization: theory and applications in multidimensional data analysis. *Recent Advances in Descriptive Multivariate Analysis*, pages 157–189.

- Hocking, T., Vert, J.-P., Bach, F., and Joulin, A. (2011). Clusterpath: an algorithm for clustering using convex fusion penalties. In L. Getoor and T. Scheffer, editors, *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, ICML '11, pages 745–752, New York, NY, USA. ACM.
- Hopcroft, J. and Tarjan, R. (1973). Algorithm 447: Efficient algorithms for graph manipulation. *Communications of the ACM*, **16**(6), 372–378.
- Hunter, D. R. and Lange, K. (2004). A tutorial on MM algorithms. *American Statistician*, **58**, 30–37.
- Kittles, R. A. and Weiss, K. M. (2003). Race, ancestry, and genes: implications for defining disease risk. *Annual Rev Genomics Hum Genet*, **4**, 33–67.
- Lange, K. (2012). *Optimization*. Springer Texts in Statistics. Springer-Verlag, second edition.
- Lange, K. and Keys, K. L. (2014). The MM proximal distance algorithm. *Proceedings 2014 International Congress of Mathematicians*, (in press).
- Lange, K., Hunter, D. R., and Yang, I. (2000). Optimization transfer using surrogate objective functions. *Journal of Computational and Graphical Statistics*, **9**, 1–20.
- Li, J. Z., Absher, D. M., Tang, H., Southwick, A. M., Casto, A. M., Ramachandran, S., Cann, H. M., Barsh, G. S., Feldman, M., Cavalli-Sforza, L. L., and Myers, R. M. (2008). Worldwide human relationships inferred from genome-wide patterns of variation. *Science*, **319**(5866), 1100–1104.
- Lindsten, F., Ohlsson, H., and Ljung, L. (2011). Clustering using sum-of-norms regularization: With application to particle filter output computation. In *Statistical Signal Processing Workshop (SSP), 2011 IEEE*, pages 201–204. IEEE.
- Nelson, M. R., Bryc, K., King, K. S., Indap, A., Boyko, A. R., Novembre, J., Briley, L. P., Maruyama, Y., Waterworth, D. M., Waeber, G., Vollenweider, P., Oksenberg, J. R., Hauser, S. L., Stirnadel, H. A., Kooner, J. S., Chambers, J. C., Jones, B., Mooser, V., Bustamante, C. D., Roses, A. D., Burns, D. K., Ehm, M. G., and Lai, E. H. (2008). The Population Reference Sample, POPRES: a resource for population, disease, and pharmacological genetics research. *Am. J. Hum. Genet.*, **83**(3), 347–358.
- Parikh, N. and Boyd, S. (2013). Proximal algorithms. *Foundations and Trends in Optimization*, **1**(3), 123–231.
- Perou, C. M., S7rlie, T., Eisen, M. B., van de Rijn, M., Jeffrey, S. S., Rees, C. A., Pollack, J. R., Ross, D. T., Johnsen, H., Akslen, L. A., Fluge, O., Pergamenschikov, A., Williams, C., Zhu, S. X., L7nning, P. E., B7rresen-Dale, A. L., Brown, P. O., and Botstein, D. (2000). Molecular portraits of human breast tumours. *Nature*, **406**(6797), 747–752.
- Price, A. L., Patterson, N. J., Plenge, R. M., Weinblatt, M. E., Shadick, N. A., and Reich, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.*, **38**(8), 904–909.
- Pritchard, J. K., Stephens, M., and Donnelly, P. (2000). Inference of population structure using multilocus genotype data. *Genetics*, **155**(2), 945–959.
- Qamar, R., Ayub, Q., Mohyuddin, A., Helgason, A., Mazhar, K., Mansoor, A., Zerjal, T., Tyler-Smith, C., and Mehdi, S. Q. (2002). Y-chromosomal DNA variation in Pakistan. *Am. J. Hum. Genet.*, **70**(5), 1107–1124.
- Rocheftort, H., Glondou, M., Sahla, M. E., Platet, N., and Garcia, M. (2003). How to target estrogen receptor-negative breast cancer? *Endocr. Relat. Cancer*, **10**(2), 261–266.
- Rosenberg, N. A., Pritchard, J. K., Weber, J. L., Cann, H. M., Kidd, K. K., Zhivotovsky, L. A., and Feldman, M. W. (2002). Genetic structure of human populations. *Science*, **298**(5602), 2381–2385.

- Rosenberg, N. A., Mahajan, S., Gonzalez-Quevedo, C., Blum, M. G., Nino-Rosales, L., Ninis, V., Das, P., Hegde, M., Molinari, L., Zapata, G., Weber, J. L., Belmont, J. W., and Patel, P. I. (2006). Low levels of genetic divergence across geographically and linguistically diverse populations from India. *PLoS Genet.*, **2**(12), e215.
- Sokal, R. R. and Michener, C. D. (1958). A statistical method for evaluating systematic relationships. *University of Kansas Science Bulletin*, **38**, 1409–1438.
- Su, Y.-S., Gelman, A., Hill, J., and Yajima, M. (2011). Multiple imputation with diagnostics (mi) in r: Opening windows into the black box. *Journal of Statistical Software*, **45**(2), 1–31.
- Wu, T. T. and Lange, K. (2010). The MM alternative to EM. *Statistical Science*, **25**, 492–505.

Table 1: Error rates as a function of noise in the Iris Data

Noise level $c$	HCLUST UPGMA	CONVEXCLUSTER		
		k=5	k=10	k=15
0.02	0.193	0.105	0.095	0.102
0.04	0.194	0.108	0.107	0.105
0.06	0.197	0.107	0.104	0.105
0.08	0.216	0.111	0.107	0.126
0.10	0.226	0.120	0.121	0.134

Table 2: Error rates as a function of missingness in the Iris Data

Proportion of rows with a missing attribute $c$	HCLUST UPGMA	CONVEXCLUSTER		
		k=5	k=10	k=15
0.25	0.210	0.109	0.115	0.127
0.50	0.205	0.127	0.137	0.133
0.75	0.228	0.147	0.148	0.141
1.00	0.228	0.153	0.181	0.146

Table 3: Average runtimes in seconds for different analyses

Analysis	Datapoints	Variables	CLUSTERPATH	CONVEXCLUSTER	
				CPU	GPU
HGDP	52	4,682	8.67	1.46	.32
POPRES	370	10	2.53	1.21	4.29
Breast Cancer data	16	9,216	3.14	2.37	.43

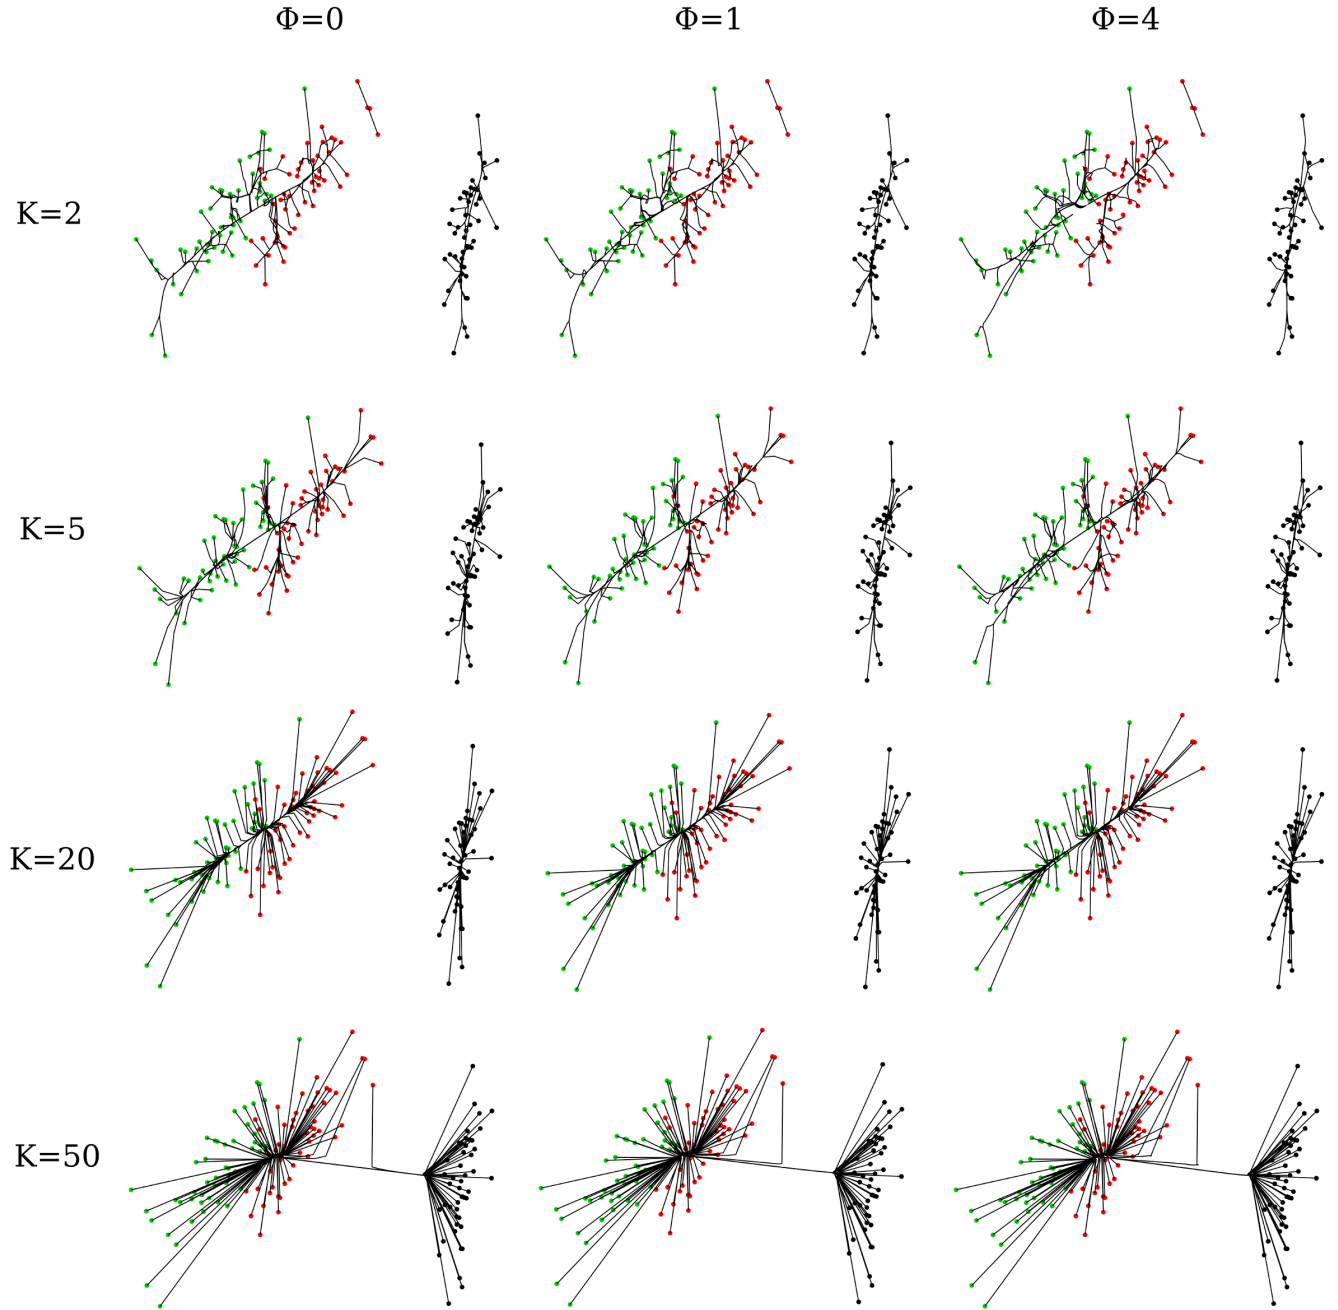


Figure 1: Convex clustering of the Iris data. Black, red, and green points denote the species *Iris-setosa*, *Iris-versicolor*, and *Iris-virginica*, respectively. These points are projections of the Iris dataset on the first two principal components (PCs). Lines trace the cluster centers as they traverse the regularization path.




Figure 2: Convex clustering of the HGDP data using a large number  $k$  of nearest neighbors to infer intercontinental connections.



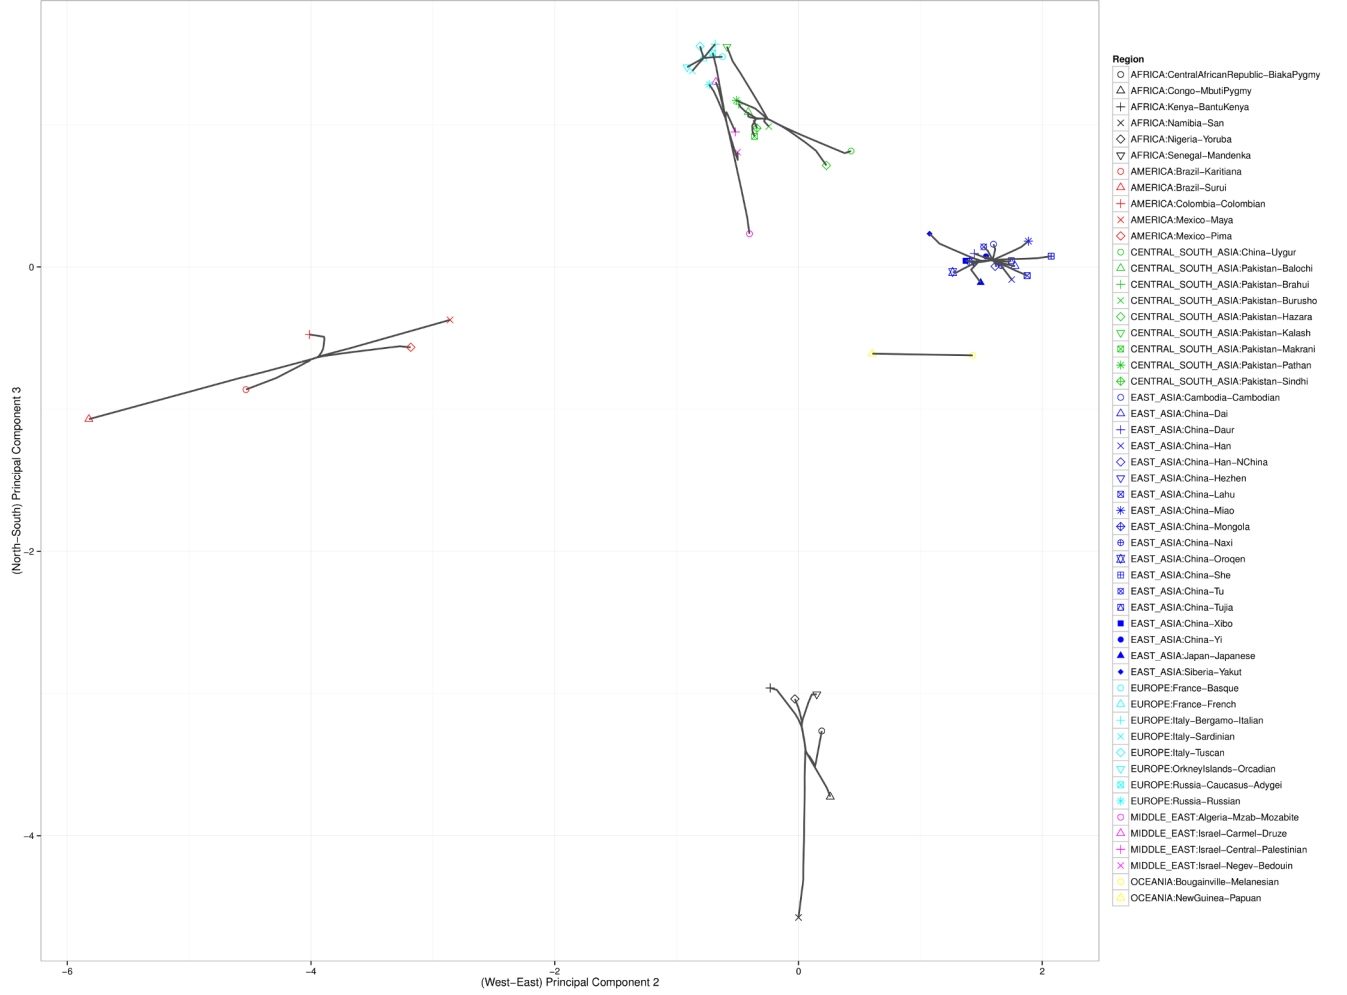


Figure 3: Convex clustering of the HGDP data using a small number  $k$  of nearest neighbors to resolve intracontinental connections.




Figure 4: Magnified view of the convex clustering results for the HGDP data in East Asia.




Figure 5: Magnified view of the convex clustering results for the HGDP data in Europe and Central Asia.




Figure 6: Hierarchical clustering of the 52 populations from the HGDP data.




Figure 7: Convex clustering of the European populations from the POPRES data using  $\phi = 0$  and  $k = 40$ .




Figure 8: Convex clustering of the European populations from the POPRES data using  $\phi = 10$  and  $k = 40$ .




Figure 9: Magnified view of results from convex clustering of Southeast Europe.




Figure 10: Magnified view of results from convex clustering of Northeast Europe.






Figure 11: UPGMA dendrogram showing genetic relationships among populations in and near the British Isles.




Figure 12: Convex clustering of the breast cancer samples. Points on the plot indicate data vectors projected onto the first and third principal components (PCs) of the sample. Lines trace the cluster centers as they traverse the regularization path.

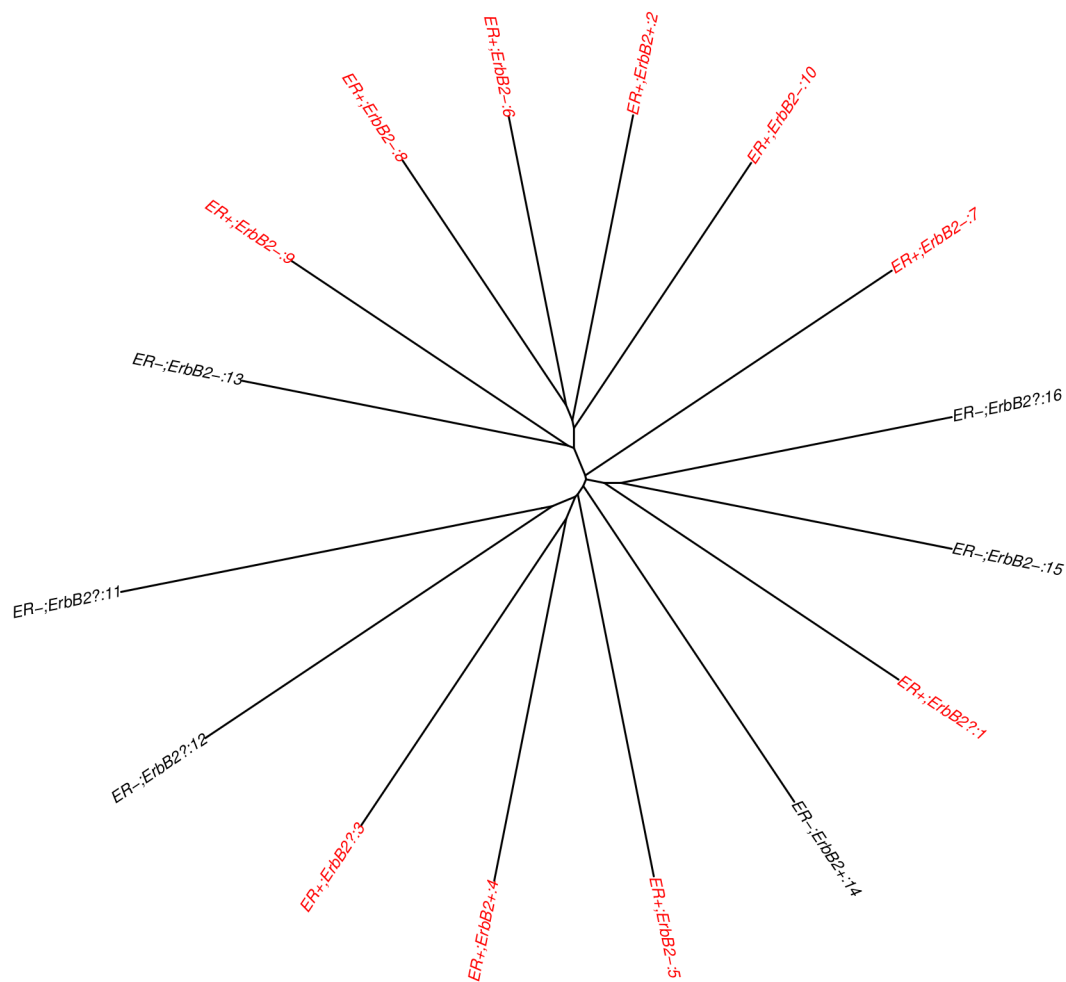


Figure 13: Average linkage hierarchical clustering of the breast cancer samples.